

embodiments shown in FIG. 2.) Micro-controller 850 is illustratively the PIC16LC774 available from Microchip. DC-to DC regulator 805 receives a DC input voltage 801 (e.g., from power supply 10 of FIG. 1) and provides a DC output voltage to device logic 895. Similarly, DC-to DC regulator 815 receives a DC input voltage 809 (e.g., from

5 power supply 10 of FIG. 1) and provides a DC output voltage to device logic 890.

Device logic 890 and device logic 895 represent the circuitry on the remainder of board 800 to which power is supplied by the DC-to-DC regulators 805 and 810. As can be observed from FIG. 8, micro-controller 850 monitors a number of power-related

parameters. In particular, the output voltage from each DC-to-DC regulator via signal

10 paths 811 and 804, temperature from each DC-to-DC regulator via signal paths 812 and 803, and temperature from two different areas of the aboard, as associated with device

logic 890 and device logic 895, via signals 891 and 894. With respect to the temperature monitoring it is assumed that each DC-to-DC regulator provides this feature and that temperature-sensing circuitry (not shown) is included within device logic 890 and device

15 logic 895. In accordance with any of the above-described flow charts of FIGs. 3, 4, 6 and 7, micro-controller 850 controls each DC-to-DC regulator via signal paths 856 and 857 and provides information to, and receives information from, communications bus 851 (e.g., the above-described system maintenance bus 7 of FIG. 1 and 5). Keep alive power is provided to micro-controller 850 via signal path 852.

20 Turning now to FIG. 9, a board 900 comprises a power control element as represented by micro-controller 950. Micro-controller 950 is illustratively the PIC16LC774 available from Microchip. Board 900 interfaces to the remainder of the system via hot plug control circuit 960. Board 900 represents a fan control board and comprises an interface to a fan as represented by fan 990, to which power is supplied via signal path 909. In addition, board 900 provides an ambient temperature sensor 975. As can be observed from FIG. 9, micro-controller 950 monitors ambient temperature sensor 975 via signal path 976 and a revolutions-per-minute (rpm) signal from fan 990, via signal path 992. In accordance with any of the above-described flow charts of FIGs. 3, 4, 6 and 7, micro-controller 950 controls fan 990 either to change the rpm value, via signal 25 path 996, and/or shutdown fan 990, via signal path 997. In addition, micro-controller 950

provides information to, and receives information from, communications bus 851 (e.g., the above-described system maintenance bus 7 of FIG. 1 and 5). Keep alive power is provided to micro-controller 950 via signal path 852.

Another server arrangement in accordance with the principles of the invention is shown in FIG. 10. As before, the elements shown in FIG. 10, other than the inventive concept, are well known and not described in detail. Server 400 comprises a number of racks of equipment (only those relevant to the invention are shown). Like numbers indicate like elements described earlier. The top rack is an array of fans controlled by, e.g., board 900 of FIG. 9. The next two racks are arrays of boards 800-1 through 800-N, all of which support distributed power control as described above, and which are represented by board 800 of FIG. 8. One rack comprises a system maintenance card 450 (described further below). Another rack comprises console 520 (i.e., the person computer functionality is integrated into server 400). Console 520 interfaces to a separate monitor, keyboard and mouse525. The last rack, represented by power supply 10, provides power to server 400.

Server 400 provides for distributed power control of boards 800-1, ..., 800-N and board 900, as described above. One variation illustrated by server 400 is the inclusion of a system maintenance card 450, which also provides the port 510 functionality described above. System maintenance card 450 serves as a master power controller for the system and controls, e.g., boards 800-1, ..., 800-N and board 900 via system maintenance bus 7 using the above-mentioned I<sup>2</sup>C signaling. System maintenance card 450 also provides the interface to console 520 via system maintenance bus 9 (also using but not limited to, illustratively, JTAG signaling). An illustrative embodiment of system maintenance card 450 is shown in FIG. 11. System maintenance card 450 comprises hot plug control circuit 460 and micro-controller 455, which receives power via signal path 451 and is coupled to system maintenance bus 7 and system maintenance bus 9. Micro-controller 455 is a stored-program based processor and memory such as the Vitesse VSC215. System maintenance card 450 provides a central interface point between console 520 and server 400. The inclusion of system maintenance card 450 provides the ability to also implement the flow charts represented by FIGs. 3, 4, 6 and 7 directly on system

maintenance card 450. Thus, system maintenance card 450 can receive information on a power-related parameter from a board, e.g., board 800-1, and, if necessary, instruct the power control element of board 800-1 to adjust the regulator, or, perhaps, shut the board down. Conversely, system maintenance card 450 passes along board-specific information

5 to console 520 for inclusion in, e.g., the above-mentioned log file, and also passes along maintenance instructions from console 520 to other parts of the server.

As described above, the inventive concept provides a scaleable power solution.

Indeed, the inventive concept provides the ability for a system, such as a multi-processor computer system, to monitor and adjust power-related parameters – on a per board basis.

10 For example, the inventive concept provides, and is not limited to, the ability to monitor voltages, voltage margins, temperature, and/ or current; the ability to predict faults; and the ability to perform fault isolation (e.g., by shutting down a board).

The foregoing merely illustrates the principles of the invention and it will thus be appreciated that those skilled in the art will be able to devise numerous alternative arrangements which, although not explicitly described herein, embody the principles of the invention and are within its spirit and scope. For example, consider the following.

15 Although the inventive concept was described in the context of a server having  $N$  boards, the server is not limited to only  $N$  boards; and it is not required that every board in the server incorporate the invention. Similarly, although the inventive concept was described in the context of a signaling path that conveyed power to a board and a signaling path that conveyed power to a power control element of the board, fewer, or more, signaling paths may be used. Also, for simplicity the distributed power control was described in the context of a voltage source – it should be realized that the inventive concept also applies to, e.g., a current source. In addition, this signaling used, e.g., I<sup>2</sup>C and JTAG, were

20 merely illustratively and other types of signaling can be used.